Multi-Agent Reinforcement Learning for Energy Harvesting Two-Hop Communications with Full Cooperation

Andrea Ortiz, *Student Member, IEEE,* Hussein Al-Shatri, *Member, IEEE,* Tobias Weber, *Senior Member, IEEE,* and Anja Klein, *Member, IEEE*

arXiv:1702.06185v1 [cs.IT] 8 Feb 2017

Abstract

We focus on energy harvesting (EH) two-hop communications since they are the essential building blocks of more complicated multi-hop networks. The scenario consists of three nodes, where an EH transmitter wants to send data to a receiver through an EH relay. The harvested energy is used exclusively for data transmission and we address the problem of how to efficiently use it. As in practical scenarios, we assume only causal knowledge at the EH nodes, i.e., in each time interval, the transmitter and the relay know their own current and past amounts of incoming energy, battery levels, data buffer levels and channel coefficients for their own transmit channels. Our goal is to find transmission policies which aim at maximizing the throughput considering that the EH nodes fully cooperate with each other to exchange their causal knowledge during a signaling phase. We model the problem as a Markov game and propose a multi-agent reinforcement learning algorithm to find the transmission policies. Furthermore, we show the trade-off between the achievable throughput and the signaling required, and provide convergence guarantees for the proposed algorithm. Results show that even when the signaling overhead is taken into account, the proposed algorithm outperforms other approaches that do not consider cooperation among the nodes.

Index Terms

Two-hop communications, energy harvesting, decode and forward, multi-agent reinforcement learning, linear function approximation.

This work was funded by the LOEWE Priority Program NICER under grant No. III L5-518/81.004.

I. Introduction

Energy harvesting (EH) wireless communications refers to scenarios where the wireless communication nodes are able to collect energy from the environment. The wireless communication nodes may be battery-operated low power devices, especially if wireless sensor networks are considered. Depending on the specific application, the replacement of their batteries can be too expensive or sometimes infeasible [1], e.g., in the cases where the nodes are located in remote locations, inside structures, or even inside the human body. In order to provide sustainable service or to reduce the operating expenses, EH has been considered as a promising energy source for such wireless communication nodes. In EH wireless communication networks, the EH capability of the nodes increases the network lifetime and can lead to perpetual operation because the nodes can use the harvested energy to recharge their batteries [2], [3]. However, the benefits of EH are not limited to an increased network lifetime. The fact that the EH nodes can collect energy from natural or man-made sources, e.g., solar, chemical or electromagnetic radiation, helps to reduce greenhouse gas emissions. Furthermore, since the EH nodes can work independently of the power grid, EH wireless communication networks can be deployed in areas that are usually hard to reach. In this paper, we address the problem of how to efficiently use the harvested energy and we tackle the problem from a communications perspective, i.e., we discuss how to efficiently transmit data using the harvested energy as the only energy source.

In an EH scenario, the communication range depends on the amount of harvested energy at the EH transmitter. This amount of harvested energy varies according to the energy source that is considered. For example, for energy harvesting based on electromagnetic radiation, the power density is in the order of fractions of $\mathrm{nW/cm^2}$, and for solar energy, is in the order of hundreds of $\mathrm{mW/cm^2}$. To increase the limited communication range in an EH communication scenario, relaying techniques can be considered since they are cost effective solutions for increasing the coverage, throughput and robustness of wireless networks [4], [5]. By using relaying techniques, the communication between a transmitter and a receiver which are located far apart can be achieved by introducing one or more intermediate relays for reducing the communication range of each hop. The reduction of the communication range implies a reduction of the amount of energy required for data transmission in each hop. We focus on the case where only a single EH relay is used to assist the communication between an EH transmitter and a receiver, i.e.,

EH two-hop communications. This scenario is the essential building block of more complicated EH multi-hop communication networks and exhibits all important challenges that need to be addressed when using relaying techniques, i.e., the design of transmission policies for the EH transmitter and the EH relay considering the amount of energy that is available to each of them. Our goal is to design transmission policies aiming at an efficient use of the harvested energy at the transmitter and at the relay.

A. Related Work

The study of EH wireless communications has been based on three different approaches, namely, offline approaches [6]–[21], online approaches [22]–[26] and learning approaches [27]–[30]. The offline approaches assume complete non-causal knowledge regarding the EH processes, the data arrival processes and the channel fading processes. Although this assumption cannot be fulfilled in reality, the offline approaches are useful to derive upper bounds of the performance. A more relaxed assumption is considered by the online approaches where only statistical knowledge about the EH processes, the data arrival processes and the channel fading processes is assumed to be available in advance. However, in real scenarios this statistical knowledge might not be available, especially if non-stationary EH, data arrival and channel fading processes are considered. The requirement of complete non-causal knowledge (offline approaches) or statistical knowledge (online approaches) regarding the EH processes, the data arrival processes and the channel fading processes can be overcome if a learning approach is considered. This is due to the fact that in learning approaches, more specifically in reinforcement learning, an agent learns how to behave in an unknown environment by interacting with it. In the case of EH communications, the agent can be the EH transmitter and the environment includes the unknown EH process, the data arrival process and the channel fading process. In the following, we give an overview of the state of the art of offline, online and learning approaches, first for EH point-to-point communications and secondly for EH two-hop communications.

Offline EH point-to-point communications have been investigated in [6]–[10]. In [6], the throughput maximization problem within a deadline is considered. It is shown that in this scenario, the throughput maximization problem within a deadline is equivalent to the minimization of the completion time given that a fixed amount of data needs to be transmitted. A similar scenario is investigated in [7], where the authors consider a fading channel between the transmitter and

the receiver. In this case, a modified water-filling algorithm, termed directional water-filling, is proposed to maximize the throughput within a deadline. In [8], the minimization of the distortion of the received messages is considered. It is assumed that each message has to be reconstructed at the destination within a certain deadline and a convex optimization problem for the power allocation is formulated. Additionally, the processing costs at the transmitter in a point-to-point scenario are analyzed in [9]. In [10], we consider the case when each data packet to be sent has an individual deadline. For this scenario, we formulate optimization problems to consider the delay-constrained throughput maximization problem as well as the delay-constrained energy minimization problem. Online approaches for point-to-point scenarios are investigated in [22] and [23]. In [22], statistical information about the distribution of the importance of the messages is assumed, and an on-off mechanism at the transmitter is considered. For each data packet that arrives at the transmitter, a binary decision on whether to transmit or drop the data packet is made. A save-then-transmit protocol for system outage minimization is considered in [23]. For each time interval, the harvested energy is modelled as a random variable. Additionally, it is assumed that a fixed amount of data needs to be transmitted during each time interval. Learning approaches have been applied to EH point-to-point scenarios in [27] and [28]. In [27], it is assumed that the amount of incoming energy, the channel coefficients, and the transmit power in each time interval are taken from a finite discrete set and the well-known reinforcement learning algorithm Q-learning is applied in order to maximize the throughput in a fixed period of time. In our previous work [28], we combine the reinforcement learning algorithm state-action-reward-state-action (SARSA) with linear function approximation to enable the use of incoming energy and channel values which are taken from a continuous range. With this approach, we are able to improve the performance in the EH point-to-point scenario as compared to the case when Q-learning is used.

For EH two-hop communications, offline approaches have been the major direction of state of the art research [15]–[19]. In [15], the throughput maximization problem within a deadline is studied and two cases are distinguished, namely a full-duplex and a half-duplex relay. For the case of a full-duplex relay, an optimal transmission scheme is provided. However, in the half-duplex case, a simplified scenario is assumed where a single energy arrival is considered at the transmitter. In [16], the authors formulate a convex optimization problem to find offline transmission policies for multiple parallel relays in a decode-and-forward EH two-hop commu-

nication scenario. Half-duplex amplify-and-forward EH two-hop communications are considered in our previous work [17]. In this case, we used D.C. programming to find the optimal power allocation. In [18], the throughput maximization problem is investigated when the transmitter harvests energy multiple times and the amplify-and-forward relay has only one energy arrival. In [19], the impact of a finite data buffer at the relay is investigated. Similar to the previous case, it is assumed that the transmitter harvests energy several times while the relay harvests energy only once. As mentioned before, the advantage of the offline approaches is that they provide an upper bound of the performance. However, the assumption of complete non-causal knowledge is hard to fulfil in real implementations. In [25] and [26], online approaches are considered. In [25], a half-duplex amplify-and-forward EH two-hop communication scenario is studied. The authors assume statistical knowledge about the energy harvesting process and find the transmission policy using discrete dynamic programming. A similar scenario is considered in [26], where the power allocation policy is found using Lyapunov optimization techniques. Despite the fact that the online approaches do not require non-causal knowledge as the offline approaches do, their application is restricted to the case of stationary processes or when the statistics of the random processes are known. The learning approaches have the advantage that they do not require a specific model or statistical knowledge of the EH processes, the data arrival processes or the channel fading processes. A learning approach for an EH two-hop communication scenario is considered in our previous work [30]. In this case, it is assumed that the transmitter and the relay are not able to exchange their causal knowledge. Consequently, the two-hop communication scenario is separated into two point-to-point scenarios and the transmitter and the relay solve independent reinforcement learning problems to find the transmission policies that aim at maximizing the throughput. In the following in this paper, we study the case when the transmitter and the relay exchange their causal knowledge during a signaling phase and exploit this knowledge to improve the achievable throughput.

B. Contributions

As previously mentioned, we focus our work on EH two-hop communications. In contrast to the state of the art, we consider a realistic scenario in which each EH node, i.e, the transmitter and the relay, have only causal knowledge. This means, in each time interval, each EH node knows its current and past amounts of incoming energy, its current and past battery levels, its

current and past data buffer levels and its current and past channel coefficients for its transmit channel. For the transmitter, this channel coefficient corresponds to the channel between the transmitter and the relay, and for the relay, this channel coefficient corresponds to the channel between the relay and the receiver. As the throughput in a two-hop communication scenario depends on the transmission policies of both, the transmitter and the relay, we consider the case when the EH nodes fully cooperate with each other to exchange their causal knowledge during a signaling phase. This cooperation provides each EH node with the causal knowledge of the other EH node. The EH nodes exploit the obtained causal knowledge to adapt their transmission policies. We are interested in a distributed solution where each EH node finds its own transmission policy taking into account its own causal knowledge and the knowledge it has obtained during the signaling phase. Consequently, we model this scenario as a Markov game because it provides the generalization of the model for decision-making situations given by Markov decision processes which considers only a single decision-making agent, to the case when multiple decision-making agents are considered [31]. Additionally, we propose a multi-agent reinforcement learning algorithm to find the transmission policies at the transmitter and at the relay.

To validate our proposed algorithm, we derive convergence guarantees based on reinforcement learning. Moreover, by numerical results we show that the performance of the proposed algorithm has only a small degradation compared to the offline case which requires complete non-causal knowledge. Additionally, we show that even when the overhead caused by the signaling phase is taken into account, the proposed algorithm outperforms other approaches that do not consider cooperation among the EH nodes and therefore, do not require a signaling phase.

C. Organization of the paper

The rest of the paper is organized as follows. In Section II, the system model is presented and the transmission scheme is explained. In Section III, the problem of EH two-hop communications with fully cooperating nodes is addressed. We model the problem as a Markov game and apply multi-agent reinforcement learning to find the transmission policies at the transmitter and at the relay. Convergence guarantees for the proposed algorithm are presented in Section IV. Numerical performance results are presented in Section V and Section VI concludes the paper.

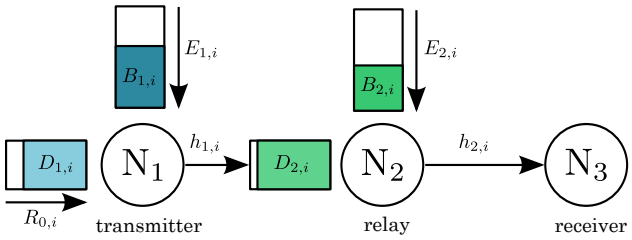

Fig. 1: EH two-hop communication scenario.

II. System Model

An EH two-hop communication scenario consisting of three single-antenna nodes is considered. The symbol N_k, $k \in \{1,2,3\}$, is used to label the nodes. As depicted in Fig. 1, the transmitter N_1 wants to transmit data to the receiver N_3. It is assumed that the link between these two nodes is weak. Therefore, the nodes cannot communicate directly. To enable the communication, node N_2 acts as a full-duplex decode-and-forward relay. It is assumed that the relay N_2 forwards the data from node N_1 to node N_3 and it is able to perfectly cancel the self-interference caused by its transmission. Furthermore, a time slotted system using I time intervals is considered with a constant duration τ for each time interval i, $i = 1, ..., I$.

The nodes N_1 and N_2 harvest energy from the environment and use this energy exclusively for data transmission. We consider a discrete time model in which at the beginning of each time interval i, an amount of energy $E_{l,i} \in \mathbb{R}^+$, $l \in \{1,2\}$ is received by node N_l. The amount of energy $E_{l,i}$ may also take the value $E_{l,i} = 0$ to include the case when node N_l does not harvest energy in time interval i. The maximum amount of energy that can be harvested at node N_l, termed $E_{\max,l}$, depends on the energy source that is used.

The amount of harvested energy $E_{l,i}$ is stored in a rechargeable finite battery with maximum capacity $B_{\max,l}$. It is assumed that no energy is lost in the process of storing or retrieving energy from the batteries. Additionally, it is assumed that the batteries cannot be recharged instantaneously. Consequently, the amount of harvested energy $E_{l,i}$ cannot be used in time interval i but earliest in time interval $i+1$. This model, is similar to the one presented in [23], where the authors assume that they have one main battery which powers the transmitter. However, since the battery cannot be recharged and discharged simultaneously, they assume there is a secondary battery which is charged throughout the time interval while the energy in the main battery is

used for data transmission. After the data transmission is over, all the energy stored in the secondary battery is passed to the main battery (which can be seen as an amount of energy that comes at the beginning of the next time interval). The battery levels $B_{l,i}$ are always measured at the beginning of each time interval i. Furthermore, it is assumed that at the beginning of time interval $i = 1$, the nodes have not yet harvested any energy and their batteries are empty, i.e., $B_{l,1} = 0, l \in \{1, 2\}$.

The data available for transmission at node N_1 depends on its own data arrival process. It is assumed that at the beginning of time interval i, a data packet of $R_{0,i}$ bits is received by node N_1 and the incoming data is stored in a finite data buffer with size $D_{\text{max},1}$, measured in bits. Moreover, it is assumed that node N_2 does not have any own data to transmit to the other nodes. Consequently, node N_2 can only retransmit what it has received from node N_1. Similar to node N_1, node N_2 receives $R_{1,i}$ bits in time interval i and stores them in its data buffer. The maximum amount of data which node N_2 can store is limited by the size of its data buffer which is given by $D_{\text{max},2}$. The data buffer level of node N_l is measured at the beginning of time interval i and is denoted by $D_{l,i}$. It is assumed that at at the beginning of time interval $i = 1$, both data buffers are empty, i.e., $D_{l,1} = 0$.

The fading channel from node N_1 to node N_2 is described by the channel coefficient $h_{1,i} \in \mathbb{C}$ while the fading channel between node N_2 and node N_3 is described by the channel coefficient $h_{2,i} \in \mathbb{C}$. It is assumed that the channels stay constant for the duration of one time interval. The noise at nodes N_2 and N_3 is assumed to be independent and identically distributed (i.i.d.) zero mean additive white Gaussian noise with variance $\sigma_2^2 = \sigma_3^2 = \sigma^2$. Additionally, a bandwidth W is assumed to be available at the EH nodes for the transmission from node N_1 to node N_2 and from node N_2 to node N_3.

In constrast to our previous work [30], where an EH two-hop communication scenario in which the nodes do not cooperate with each other was studied, in this paper we consider that the EH nodes fully cooperate and exchange their causal knowledge with each other. The transmission scheme is illustrated in Fig. 2. To exchange the causal knowledge, a signaling phase of duration τ_{sig} is included in each time interval. After the signaling phase, the nodes transmit data during the rest of the time interval. The time available for the data transmission is $\tau_{\text{data}} = \tau - \tau_{\text{sig}}$.

The signaling phase provides each EH node with the causal knowledge of the other EH node. To exchange their causal knowledge, the EH nodes use a transmit power $p_{\text{sig},l,i}$ which is kept

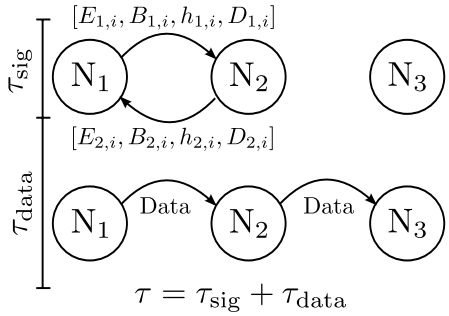

Fig. 2: Communication schemes

constant during τ_{sig}. The value of τ_{sig} depends on the tolerable quantization error, the available bandwidth and the channel coefficients, and it is explained in detail in Section III-E. When node N_l does not have enough energy to transmit during the signaling phase, i.e., $B_{l,i} < \tau_{\text{sig}} p_{\text{sig},l,i}$, $p_{\text{sig},l,i}$ is set to zero and the node does not transmit anything during τ_{sig}. In every time interval, after τ_{sig}, nodes N_1 and N_2 decide independently on the transmit power $p_{l,i}$ to be used for the transmission of data and $p_{l,i}$ is kept constant during τ_{data} [6]. The throughput achieved by node N_l in time interval i is the amount of data received by node N_{l+1} which is measured in bits and is given by

$$R_{l,i} = \tau_{\text{data}} W \log_2 \left(1 + \frac{|h_{l,i}|^2 p_{l,i}}{\sigma^2}\right). \tag{1}$$

Only the energy already stored in the battery can be used for the transmission of either the signaling or the data. As a result, the following energy causality constraint has to be fulfilled in every time interval:

$$\tau_{\text{sig}} p_{\text{sig},l,i} + \tau_{\text{data}} p_{l,i} \leq B_{l,i}. \tag{2}$$

Moreover, in the selection of $p_{l,i}$ the finite capacities of the batteries have to be considered and battery overflow situations, in which part of the harvested energy is lost because the batteries are full, should be avoided. The battery overflow constraint is given by

$$B_{l,i} + E_{l,i} - \tau_{\text{data}} p_{l,i} - \tau_{\text{sig}} p_{\text{sig},l,i} \leq B_{\text{max},l}. \tag{3}$$

In addition to the energy causality and battery overflow constraints given in (2) and (3), the data arrival process at node N_l should be considered. Only data already stored in the data buffer can be transmitted. Therefore, the data causality constraint

$$R_{l,i} \leq D_{l,i} \tag{4}$$

has to be fulfilled in every time interval. Additionally, in order to avoid losing data, data buffer overflow situations should be avoided. However, it should be noted that data buffer overflow situations cannot always be avoided because the transmission of data depends on the available energy. As we aim at reducing the number of data buffer overflow situations in order to maximize the throughput, we define the data buffer overflow condition in an analogous way to the battery overflow constraint in (3) as

$$D_{l,i} + R_{l-1,i} - R_{l,i} \leq D_{\max,l}. \tag{5}$$

III. Proposed Multi-Agent Reinforcement Learning for EH Two-Hop Communications with Full Cooperation

In this section, we model the EH two-hop communication problem as a Markov game and then introduce the proposed multi-agent reinforcement learning algorithm. The proposed algorithm is used to find the transmission policies at the transmitter and at the relay aiming at maximizing the throughput. For this purpose, we propose to include a signaling phase in which the nodes exchange their causal knowledge and exploit it to adjust their transmission policies.

A. Markov game

In our scenario, the transmitter and the relay independently decide on the transmit power to use for data transmission with the aim of maximizing the achieved throughput. When only one node is considered, e.g., in a point-to-point scenario, these decision-making situations can be modelled as Markov decision processes. However, in our case, the achieved throughput depends on the transmission policies of both the transmitter and the relay. Consequently, Markov decision processes are no longer suitable because more than one node has to be considered. Markov games are a generalization of Markov decision processes and are used to model decision-making situations in which more than one agent is involved [31].

A Markov game composed of n players is defined by the set $\mathcal{S}$ of states in which the players can be, the sets $\mathcal{A}_1, ..., \mathcal{A}_n$ of actions of each player, the transition function $\mathcal{T}$ and the reward functions $\mathcal{R}_1, ..., \mathcal{R}_n$ for each player [32]. In our case, the players correspond to the transmitter and the relay. Consequently, we have to consider $n = 2$ players. Each state contains the values corresponding to both nodes, i.e., the amounts of incoming energy $E_{l,i}$ at both nodes, both battery levels $B_{l,i}$, both channel coefficients $h_{l,i}$, and both data buffer levels $D_{l,i}$. This means that in

time interval i, the corresponding state $S_i \in \mathcal{S}$ is the tuple [$E_{1,i}$, $E_{2,i}$, $B_{1,i}$, $B_{2,i}$, $h_{1,i}$, $h_{2,i}$, $D_{1,i}$, $D_{2,i}$]. The set $\mathcal{S}$ comprises an infinite number of states S_i because the amounts of incoming energy, the battery levels, the channel coefficients and the data buffer levels can take any value in a continuous range. As in our previous work [28], [30], we define the sets of actions $\mathcal{A}_1$ and $\mathcal{A}_2$ for nodes N_1 and N_2, respectively, as finite sets in order to simplify the selection of the transmit powers values. The sets of actions are defined as $p_{l,i} \in \mathcal{A}_l = \{0, \delta, 2\delta, ..., B_{\text{max},l}\}$, where δ is a step size. The transition function $\mathcal{T}$ is defined as $\mathcal{T} : \mathcal{S} \times \mathcal{A}_1 \times \mathcal{A}_2 \rightarrow \mathcal{S}$ and it specifies that, given state S_i and by selecting transmit powers $p_{1,i} \in \mathcal{A}_1$ and $p_{2,i} \in \mathcal{A}_2$, the nodes reach state S_{i+1}, i.e., $S_{i+1} = \mathcal{T}(S_i, p_{1,i}, p_{2,i})$. The reward function $\mathcal{R}_l$ gives the immediate reward obtained by node N_l when transmit power $p_{l,i}$ is selected while being in state S_i. In our case, the nodes aim at maximizing the throughput, i.e., the amount of data received by node N_3. Consequently, nodes N_1 and N_2 share the same objective $\mathcal{R}_1 = \mathcal{R}_2 = \mathcal{R}$. Moreover, the reward $R_i \in \mathcal{R}$ corresponds to the throughput achieved in time interval i for $l = 2$ and it is calculated using (1).

As node N_l has only causal knowledge about its state, it does not know how much energy will be harvested, how much data will arrive or how the channel will be in future time intervals. As in our previous work [30], we consider this uncertainty by defining the discount factor of future rewards γ, $0 \leq \gamma \leq 1$, which quantifies the preference of achieving a larger throughput in the current time interval over future ones. Our goal is to select the transmit powers $p_{l,i}$, $\forall l, i$, in order to maximize the expected throughput which is given by

$$R = \lim_{I \rightarrow \infty} \mathbb{E}\left[\sum_{i=1}^{I} \gamma^i R_{2,i}\right]. \tag{6}$$

To maximize the expected throughput, we need to find the transmission policies for nodes N_1 and N_2 which correspond to the transmit powers to be used for data transmission in each time interval. Each transmission policy π_l, $l \in \{1, 2\}$ is a mapping from a given state S_i to the transmit power $p_{l,i}$ that should be selected by each node, i.e. $p_{l,i} = \pi_l(S_{l,i})$.

B. Single-agent reinforcement learning

To facilitate the description of the proposed multi-agent reinforcement learning algorithm, let us first consider the single-agent case. For this purpose, let us assume that a central entity has, in each time interval, perfect knowledge about the state S_i and uses reinforcement learning to find

the combined policy Π, with $\Pi = (\pi_1, \pi_2)$. Π can be evaluated using the so-called action-value function $\mathrm{Q}^{\Pi}(S_i, P_i)$, with $P_i = (p_{1,i}, p_{2,i})$, which is defined as the expected reward starting from state S_i, selecting transmit power P_i and following Π thereafter [33]. The optimal policy Π^* is the policy whose action-value function's value is greater than or equal to the value obtained by any other policy for every state S_i and action P_i. The corresponding action-value function for the optimal policy Π^* is denoted by Q^*. Furthermore, determining Π^* is straightforward when Q^* is known because for each state S_i, any transmit power P_i that maximizes $\mathrm{Q}^*(S_i, P_i)$ is an optimal action.

The action-value function cannot be calculated in advance because only causal information is available at the nodes and the statistics of the EH processes, data arrival processes and channel fading processes are unknown. To overcome this, reinforcement learning builds an estimate of the action-value function Q^{Π}. Specifically, we consider the temporal-difference reinforcement learning algorithm State-Action-Reward-State-Action (SARSA) which builds the estimate based on the visited states and the obtained rewards. The SARSA update rule for the estimate of the action value function $\mathrm{Q}^{\Pi}(S_i, P_i)$ is given by

$$\mathrm{Q}^{\Pi}_{i+1}(S_i, P_i) = \mathrm{Q}^{\Pi}_{i}(S_i, P_i)(1 - \alpha_i) + \alpha_i \left[R_i + \gamma \mathrm{Q}_i(S_{i+1}, P_{i+1})\right] \tag{7}$$

[33], where α_i is a small positive fraction which influences the learning rate. Additionally, we include the sub-index i in the definition of the update to emphasize the fact that the update changes the estimated value of the action-value function for the pair (S_i, P_i) in time interval i.

C. Multi-agent reinforcement learning

We are interested in a distributed solution for finding the transmission policies for nodes N_1 and N_2. Therefore, we propose a multi-agent reinforcement learning algorithm based on the SARSA update shown in (7). Similar to [34], we consider that the nodes have a common objective which is in our case to maximize the expected throughput, and in every time interval will make decisions that aim at maximizing the throughput. As the nodes only exchange their causal knowledge and decide on their own transmit power independently, they do not know in advance the transmit power which will be selected by the other node. Consequently, they cannot build an estimate of the centralized action-value function $\mathrm{Q}^{\Pi}(S_i, P_i)$. Instead of the action-value function $\mathrm{Q}^{\Pi}(S_i, P_i)$, in the proposed multi-agent reinforcement learning algorithm, each node

builds an estimate of its local action-value function $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$. To select the transmit power $p_{l,i}$, each nodes follows the ϵ-greedy policy [33]. In the ϵ-greedy policy, each node N_l acts greedily with respect to its action-value function with a probability of $1-\epsilon$, this means

$$\Pr\left[p_{l,i} = \max_{p_{l,k}\in\mathcal{A}_l} \hat{\mathrm{q}}_l^{\pi_l}(S_i, p_{l,k})\right] = 1-\epsilon, \quad 0<\epsilon<1. \tag{8}$$

However, with a probability of ϵ, node N_l will randomly select a transmit power value from the set $\mathcal{A}_l$. This method provides a trade-off between the exploration of new transmit power values and the exploitation of the known ones [33], [35].

As explained in [34], $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ is a projection of the centralized $\mathrm{Q}^{\Pi}(S_i, P_i)$ in which the nodes will only update their current estimate of $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ if the value of the update is larger than the current one. The relation between $\mathrm{Q}^{\Pi}(S_i, P_i)$ and $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ is presented in detail in Section IV and the proposed update rule for $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ is given by

$$\mathrm{q}_{l,i+1}^{\pi_l}(S_i, p_{l,i}) = \max\left\{\mathrm{q}_{l,i}^{\pi_l}(S_i, p_{l,i}),\, \mathrm{q}_{l,i}^{\pi_l}(S_i, p_{l,i})(1-\alpha_i) + \alpha_i\left[R_i + \gamma \mathrm{q}_{l,i}^{\pi_l}(S_{i+1}, p_{l,i+1})\right]\right\}. \tag{9}$$

The action-value function $\mathrm{q}_l^{\pi_l}$ is a table in which the number of fields is equal to the number of states multiplied by the number of actions. However, in our case the number of states is infinite. This means, infinitely many values of $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ would need to be stored. Since a table with an infinite number of fields cannot be constructed, linear function approximation is used to handle the infinite number of states. With linear function approximation, $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ is represented as the linear combination of a set of M feature functions $\mathrm{f}_{l,m}(S_i, p_{l,i}),\ m = 1, ..., M$ which map the state-action pair $(S_i, p_{l,i})$ onto a feature value. The proposed feature functions are explained in the next subsection. For a given pair $(S_i, p_{l,i})$, the feature values are collected in the vector $\mathbf{f}_l \in \mathbb{R}^{M\times 1}$ and the contribution of each feature is included in the vector of weights $\mathbf{w}_l \in \mathbb{R}^{M\times 1}$. The action-value function is approximated as

$$\mathrm{q}_l^{\pi_l}(S_i, p_{l,i}) \approx \hat{\mathrm{q}}_l^{\pi_l}(S_i, p_{l,i}, \mathbf{w}_l) = \mathbf{f}_l^{\mathrm{T}}\mathbf{w}_l. \tag{10}$$

When SARSA with linear function approximation is applied, the updates are performed on the weights because they control the contribution of each feature function on $\hat{\mathrm{q}}_l^{\pi_l}(S_i, p_{l,i})$. In every time interval, the vector $\mathbf{w}_l$ is adjusted in the direction that reduces the error between $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ and $\hat{\mathrm{q}}_l^{\pi_l}(S_i, p_{l,i}, \mathbf{w}_l)$ following the gradient descent approach presented in [33]. Considering the update for $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ given in (9), we propose to update $\mathbf{w}_l$ as

$$\mathbf{w}_{l,i+1} = \mathbf{w}_{l,i} + \max\left\{0, \alpha_i\left[R_i + \gamma\, \mathbf{f}_l^{T}(S_{i+1}, p_{l,i+1})\mathbf{w}_{l,i} - \mathbf{f}_l^{\mathrm{T}}(S_i, p_{l,i})\mathbf{w}_{l,i}\right]\mathbf{f}_l(S_i, p_{l,1})\right\}. \tag{11}$$

D. Feature functions

The feature functions are generally defined based on the natural attributes of the problem which are in this case the EH processes at the EH nodes, the finite batteries, the data arrival processes, the finite data buffers and the channel fading processes. For the proposed multi-agent reinforcement learning algorithm, we consider $M = 6$ binary feature functions and set $\mathbf{f}_1 = \mathbf{f}_2 = \mathbf{f}$. The first five feature functions were proposed in our previous work [28], [30] for the EH point-to-point communication scenario and the two-hop communication scenario without cooperation among the EH nodes, respectively.

The first feature function $\mathrm{f}_1(S_i, p_{l,i})$ considers the energy causality and battery overflow constraints in (2) and (3). It indicates whether in state S_i, a given transmit power $p_{l,i}$ avoids a battery overflow situation. Additionally, it indicates whether the transmit power $p_{l,i}$ fulfils the energy causality constraint and it is written as

$$\mathrm{f}_1(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } (B_{l,i} + E_{l,i} - \tau_{\text{data}} p_{l,i} \leq B_{\max,l}) \wedge (\tau_{\text{data}} p_{l,i} \leq B_{l,i}) \\ 0, & \text{else}, \end{cases} \tag{12}$$

[28] where $\wedge$ is the logical conjunction operation.

The second feature function $\mathrm{f}_2(S_i, p_{l,i})$ considers the power allocation problem by performing water-filling between the current channel realization and an estimate of the mean value $\bar{h}_{l,i}$ of past channel realizations. The water level $\upsilon_{l,i}$ is calculated as

$$\upsilon_{l,i} = \frac{1}{2}\left(\frac{B_{l,i}}{\tau_{\text{data}}} + \frac{E_{l,i}}{\tau_{\text{data}}} + \sigma^2\left(\frac{1}{|\bar{h}_{l,i}|} + \frac{1}{|h_{l,i}|}\right)\right), \tag{13}$$

and the power allocation given by the water-filling algorithm is given by

$$p_{l,i}^{\text{WF}} = \min\left\{\frac{B_{l,i}}{\tau_{\text{data}}}, \max\left\{0, \upsilon_{l,i} - \frac{\sigma^2}{|h_{l,i}|}\right\}\right\}. \tag{14}$$

As the transmit power $p_{l,i}$ can only be selected from the discrete set $\mathcal{A}_l$, the calculated $p_{l,i}^{\text{WF}}$ is rounded such that $p_{l,i}^{\text{WF}} \in \mathcal{A}_l$ holds. $\mathrm{f}_2(S_{l,i}, p_{l,i})$ is written in [28] as

$$\mathrm{f}_2(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } p_{l,i}^{\text{WF}} = p_{l,i} \\ 0, & \text{else}. \end{cases} \tag{15}$$

The third feature function $\mathrm{f}_3(S_i, p_{l,i})$ handles the case when the harvested energy is larger than the battery capacity, i.e., $E_{l,i} \geq B_{\max,l}$. In this case, the battery should be depleted in order to

minimize the amount of energy that is lost due to overflow. $\mathrm{f}_3(S_i, p_{l,i})$ is given in [28] by

$$\mathrm{f}_3(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } (E_{l,i} \geq B_{\max,l}) \wedge \left(p_{l,i} = \delta \lfloor \frac{B_{l,i}}{\tau_{\text{data}}\delta} \rfloor\right) \\ 0, & \text{else.} \end{cases} \tag{16}$$

The fourth feature function $\mathrm{f}_4(S_i, p_{l,i})$ considers the data causality constraint in (4). Let us define $R_{l,i}^{(p_{l,i})}$ as the throughput that would be achieved if the transmit power $p_{l,i}$ is selected. $\mathrm{f}_4(S_i, p_{l,i})$ indicates if $R_{l,i}^{(p_{l,i})}$ fulfils the constraint in (4) given that there is enough energy in the battery to select it. $\mathrm{f}_4(S_i, p_{l,i})$ is calculated in [30] as

$$\mathrm{f}_4(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } \left(R_{l,i}^{(p_{l,i})} \leq D_{l,i}\right) \wedge (B_{l,i} \geq \tau_{\text{data}} p_{l,i}) \\ 0, & \text{else.} \end{cases} \tag{17}$$

The fifth feature function $\mathrm{f}_5(S_i, p_{l,i})$ aims at the depletion of the data buffers as a preventive measure against data buffer overflow situations and it is defined in [30] as

$$\mathrm{f}_5(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } p_{l,i} = \underset{p'_{l,i} \in \mathcal{A}_l}{\operatorname{argmin}} \; D_{l,i} - R_{l,i}^{(p'_{l,i})} \\ 0, & \text{else.} \end{cases} \tag{18}$$

The sixth feature function $\mathrm{f}_6(S_i, p_{l,i})$ takes the available information node N_l has about node N_j, $l, j \in \{1, 2\}$, $l \neq j$ into account and it is used to further avoid data buffer overflow situations at node N_2. We focus on the data buffer overflow of node N_2 because the data buffer level $D_{2,i}$ depends on the throughput of nodes N_1 and N_2. On the contrary, the data buffer level at node N_1 depends only on the throughput of node N_1 and its data arrival process which we cannot control. To avoid data buffer overflow situations at node N_2, each node N_l determines an estimate of the power $\bar{p}_{j,i}$ to be selected by node N_j, $l \neq j$ using the water-filling procedure described by the third feature function. With $\bar{p}_{j,i}$, the corresponding throughput $R_{j,i}^{(\bar{p}_{j,i})}$ is calculated and it is compared to the data buffer level $D_{j,i}$. If $R_{j,i}^{(\bar{p}_{j,i})} > D_{j,i}$, then $\bar{p}_{j,i}$ is scaled down to the minimum power value $\bar{p}_{j,i} \in \mathcal{A}_j$ that can be used to deplete the data buffer at node N_j. The feature function is then defined for $l = 1$ as

$$\mathrm{f}_6(S_i, p_{l,i}) = \begin{cases} 1, & \text{if } \left(R_{l,i}^{(p_{l,i})} + D_{2,i} - R_{j,i}^{(\bar{p}_{j,i})} \leq D_{\max,2}\right) \wedge \left(R_{l,i}^{(p_{l,i})} + D_{2,i} - R_{j,i}^{(\bar{p}_{j,i})} \geq 0\right) \\ 0, & \text{else.} \end{cases} \tag{19}$$

In the case $l = 2$, the indices l and j should be interchanged.

E. Signaling

As mentioned in Section II, we consider a transmission scheme which consists of a signaling phase and a data transmission phase. During the signaling phase of duration τ_{sig}, the EH nodes cooperate with each other and exchange their causal knowledge. During the data transmission phase of duration τ_{data}, the EH nodes transmit the data stored in their data buffers. To facilitate the coordination among the nodes, we keep τ_{sig} fixed and in each time interval i, we calculate the power $p_{\text{sig},l,i}$ required for the transmission of the signaling. In the following, we describe how how to compute $p_{\text{sig},l,i}$.

During τ_{sig}, nodes N_l, $l \in \{1,2\}$ transmit their current amounts of incoming energy $E_{l,i}$, their current battery levels $B_{l,i}$, their current channel coefficients $h_{l,i}$ and their current data buffer levels $D_{l,i}$. Let $x_{l,i}$ be a variable that represents any of these quantities for node N_l, i.e., $x_{l,i} \in \{E_{l,i}, B_{l,i}, h_{l,i}, D_{l,i}\}$. Then, the number $L_{x_{l,i}}$ of bits required for the transmission of each quantity $x_{l,i}$ depends on the type of quantizer that is used. For simplicity, in this paper we consider a uniform quantizer. Consequently, $x_{l,i}$ depends on the tolerable quantization error $e_{x_{l,i},\text{quant}}$ and the maximum value $V_{x_{l,i},\text{max}}$ and the minimum value $V_{x_{l,i},\text{min}}$ each of them can take. $L_{x_{l,i}}$ is calculated as

$$L_{x_{l,i}} = \left\lceil \log_2 \left(\frac{V_{x_{l,i},\text{max}} - V_{x_{l,i},\text{min}}}{e_{x_{l,i},\text{quant}}} \right) - 1 \right\rceil, \tag{20}$$

where $\lceil \cdot \rceil$ is the rounding operation to the next integer value greater than or equal to the evaluated number. For example, the number of bits required to transmit the battery level of node N_1 during the signaling phase of time interval i is denoted by $L_{B_{1,i}}$ and it is calculated using (20) with $V_{x_{l,i},\text{max}} = B_{\text{max}}$ and $V_{x_{l,i},\text{min}} = 0$ for a certain quantization error $e_{B_{1,i},\text{quant}}$. Since $V_{x_{l,i},\text{max}}$ and $V_{x_{l,i},\text{min}}$ are fixed for each $x_{l,i}$, the number of bits node N_l requires for signaling to node N_j, $l,j \in \{1,2\}$, $l \neq j$, is constant for all the time intervals and it is given by

$$L_l = \sum_{\forall x_{l,i}} L_{x_{l,i}}. \tag{21}$$

Given L_l, the power $p_{\text{sig},l,i}$ required to transmit the signaling from node N_l to node N_j is

$$p_{\text{sig},l,i} = \frac{\sigma^2}{|h_{l,i}|^2} \left(2^{\frac{L_l}{W \tau_{\text{sig}}}} - 1 \right). \tag{22}$$

It should be noted that the amount of energy $\tau_{\text{sig}} p_{\text{sig},l,i}$, used by each node for the transmission of the causal knowledge during the signaling phase is deducted from the battery level $B_{l,i}$ and

the rest is available for data transmission. Moreover, if node N_l does not have enough energy in its battery for the transmission of the signaling, it transmits nothing during τ_{sig}. However, in this case when node N_l does not have enough energy in the battery to transmit the signaling, it can still use the energy in the battery for data transmission during τ_{data}. When a node N_l cannot transmit the signaling, the other node N_j, $j \in \{1, 2\}$, $j \neq l$, assumes that no energy was harvested by node N_l, i.e., $E_{l,i} = 0$, and that the signaling was not sent because the battery level of node N_l is zero, i.e., $B_{l,i} = 0$. Additionally, since there is no knowledge about the channel coefficient, it is assumed that $h_{l,i} = h_{l,i-1}$ and for the data buffer level of node N_l, it is assumed that $D_{l,i} = \max\{0, D_{l,i-1} - R_{l,i-1}\}$, where $R_{l,i}$ is the number of bits transmitted by node N_l in time interval i. The overhead caused by the transmission of the causal knowledge during the signaling phase is measured in bits. As described in (1), this overhead is not included in the calculation of the achieved throughput.

F. Summary

The proposed multi-agent reinforcement learning algorithm for EH two-hop communication scenario is summarized in Algorithm 1. At the beginning, each node N_l initializes the values for the discount factor γ, the learning rate α, and the probability ϵ. Then, the EH nodes exchange their local causal knowledge during τ_{sig} and observe the state S_i. According to the state S_i and using the $\epsilon-$greedy policy, each node independently selects its own transmit power $p_{l,i}$ for the transmission of data. After the data transmission phase, the nodes calculate the reward which is the amount of data received by N_3, exchange their causal knowledge during a new signaling phase and observe the new state S_{i+1}. Each node selects the new transmit power $p_{l,i+1}$ using the $\epsilon-$greedy policy and updates its weights $\mathbf{w}_l$. The same procedure is repeated as long as the nodes remain operational.

IV. Convergence Guarantees

In this section, we provide convergence guarantees for the proposed multi-agent reinforcement learning algorithm. Inspired by the work of [34], we first prove that the local action-value function $\mathrm{q}_l^{\pi_l}(S_i, p_{l,i})$ is a projection of the centralized action value function $\mathrm{Q}^{\Pi}(S_i, P_i)$ that leads to the selection of the best action in $\mathrm{Q}^{\Pi}(S_i, P_i)$ for node N_l. As this proof only holds for the case of a finite number of states, we then show that the proof given in [36] for the update of weights

Algorithm 1 Multi-agent reinforcement learning algorithm

initialize γ, α, ϵ and $\mathbf{w}_l$

exchange causal knowledge and observe state S_i

select transmit power $p_{l,i}$ using the ϵ-greedy policy

while node N_l is harvesting energy **do**

transmit using the selected transmit power $p_{l,i}$

calculate corresponding reward $R_{2,i}$ ▷ Eq. (1)

exchange causal knowledge and observe state S_{i+1}

select next transmit power $p_{l,i+1}$ using the ϵ-greedy policy ▷ Eq. (8)

update $\mathbf{w}_l$ ▷ Eq. (11)

set $S_i = S_{i+1}$ and $p_{l,i} = p_{l,i+1}$

end while

in the single-agent reinforcement learning algorithm SARSA with linear function approximation holds for the proposed update rule given in (11) for the multi-agent case. [36] shows that in SARSA with linear function approximation, the weights converge to a bounded region.

Proposition 1. For an n-player Markov game defined by the tuple $(\mathcal{S}, \mathcal{A}_1, ..., \mathcal{A}_n, \mathcal{T}, \mathcal{R}_1, ..., \mathcal{R}_n)$ in which the nodes have the same reward function $\mathcal{R}_1 = ... = \mathcal{R}_n = \mathcal{R}$, $\mathcal{R} \geq 0$, the equality

$$\mathrm{q}_{l,i}(S_i, p_{l,i}) = \max_{\substack{P_i=(p_{1,i},...,p_{n,i}) \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i), \tag{23}$$

where $\mathrm{Q}_i(S_i, P_i)$ and $\mathrm{q}_{l,i}(S_i, p_{l,i})$ are the values of the centralized and local action-value function in time interval i, respectively, holds for any player l, any state S_i, and any individual action $p_{l,i}$ in the i^{th} time interval. The values of $\mathrm{Q}_i(S_i, P_i)$ and $\mathrm{q}_{l,i}(S_i, p_{l,i})$ are updated in each time interval using (7) and (9), respectively. Additionally, $P_i^{(l)}$ is the l^{th} element in P_i which corresponds to the action of player l in time interval i according to the centralized policy Π.

Proof: As in [34], the proof is done by induction on i. At $i = 1$, no reward has been obtained. Therefore, Q and q_l are zero for every state $S_1 \in \mathcal{S}$ and $p_{l,1} \in \mathcal{A}_l$, $l \in \{1, ..., n\}$ and (23) holds. For arbitrary i, (23) holds for any pair $(S_j, p_{k,j})$, $S_j \neq S_i$, $p_{k,j} \neq p_{l,i}$ because the updates in (7) and (9) are only performed on the particular pair $(S_i, p_{l,i})$. Now, to prove (23) for

the pair $(S_i, p_{l,i})$, we include the right side of (23) in the update of $\mathrm{q}_{l,i}(S_i, p_{l,i})$ in (9) as

$$\mathrm{q}_{l,i+1}(S_i, p_{l,i}) = \max \Bigg\{ \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i),\ (1-\alpha_i) \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i) + \alpha_i \left[R_i + \gamma \max_{\substack{P_{i+1} \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_{i+1}, P_{i+1}) \right] \Bigg\}. \tag{24}$$

Given the transition function $\mathcal{T}$, we have $S_{i+1} = \mathcal{T}(S_i, P_i)$. Moreover, the next action P_{i+1} depends on the policy Π, i.e., $P_{i+1} = \Pi(S_{i+1})$. As a result, (24) can be expressed as

$$\mathrm{q}_{l,i+1}(S_i, p_{l,i}) = \max \Bigg\{ \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i),\ (1-\alpha_i) \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i) + \alpha_i \left[R_i + \gamma \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i\Big(\mathcal{T}(S_i, P_i), \Pi(\mathcal{T}(S_i, P_i))\Big) \right] \Bigg\} \tag{25}$$

and by reorganizing the terms we obtain

$$\mathrm{q}_{l,i+1}(S_i, p_{l,i}) = \max \Bigg\{ \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_i(S_i, P_i),\ \max_{\substack{P_i \\ P_i^{(l)}=p_{l,i}}} \Big\{ (1-\alpha_i)\, \mathrm{Q}_i(S_i, P_i) + \alpha_i \Big[R_i + \gamma \mathrm{Q}_i\big(\mathcal{T}(S_i, P_i), \Pi(\mathcal{T}(S_i, P_i))\big) \Big] \Big\} \Bigg\}. \tag{26}$$

The second term on the right side of (26) corresponds to the update for the centralized action-value function, i.e., $\mathrm{Q}_{i+1}(S_i, P_i)$. So we can rewrite (26) as

$$\mathrm{q}_{l,i+1}(S_i, p_{l,i}) = \max \Big\{ \Big\{ \mathrm{Q}_i(S_i, P_i) \mid P_i^{(l)} = p_{l,i} \Big\} \cup \Big\{ \mathrm{Q}_{i+1}(S_i, P_i) \mid P_i^{(l)} = p_{l,i} \Big\} \Big\}. \tag{27}$$

By expanding the two terms on the right side of (27) we obtain

$$\begin{aligned} \mathrm{q}_{l,i+1}(S_i, p_{l,i}) = \max \Big\{ & \Big\{ \mathrm{Q}_i(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j \neq P_i \Big\} \cup \\ & \Big\{ \mathrm{Q}_i(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j = P_i \Big\} \cup \\ & \Big\{ \mathrm{Q}_{i+1}(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j \neq P_i \Big\} \cup \\ & \Big\{ \mathrm{Q}_{i+1}(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j = P_i \Big\} \Big\}. \end{aligned} \tag{28}$$

The first term on the right side of (28) is equal to $\mathrm{Q}_{i+1}(S_i, P_i)$ because for $P_j \neq P_i$ there is no update. The second term is always smaller than or equal to $\mathrm{Q}_{i+1}(S_i, P_i)$ because Q is monotonically increasing. Furthermore, the first and the third term are equal. Therefore, only one of them is considered. $\mathrm{q}_{l,i}(S_i, p_{l,i})$ is then written as

$$\begin{aligned}\mathrm{q}_{l,i+1}(S_i, p_{l,i}) =& \max\Big\{\Big\{\mathrm{Q}_{i+1}(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j \neq P_i\Big\} \cup \\ & \Big\{\mathrm{Q}_{i+1}(S_i, P_i) \mid P_i^{(l)} = p_{l,i}, P_j = P_i\Big\}\Big\} \\ =& \max_{\substack{P_i=(p_{1,i},\ldots,p_{n,i}) \\ P_i^{(l)}=p_{l,i}}} \mathrm{Q}_{i+1}(S_i, P_i). \end{aligned} \tag{29}$$

■

As mentioned before, the previous proof only applies to the case when a finite number of states is considered. For an infinite number of states, we use linear function approximation and propose an update rule of the weights of the approximation in (11). In [36], the author proves that single-agent SARSA with linear function approximation converges to a region. It is straightforward to see that the same proof applies to the proposed update in (11). The reasoning is as follows. From (11), it is clear that the weights are only updated if the condition

$$\alpha_i \left[R_i + \gamma\, \mathbf{f}_l^T(S_{i+1}, p_{l,i+1})\mathbf{w}_{l,i} - \mathbf{f}_l^{\mathrm{T}}(S_i, p_{l,i})\mathbf{w}_{l,i}\right] \mathbf{f}_l(S_i, p_{l,1}) > 0 \tag{30}$$

is fulfilled. In this case, the update rule of the weights corresponds to the one for single-agent SARSA which is given by

$$\mathbf{w}_{l,i+1} = \mathbf{w}_{l,i} + \alpha_i \left[R_{l,i} + \gamma \mathbf{f}^{\mathrm{T}}(S_{l,i+1}, P_{l,i+1})\mathbf{w}_{l,i} - \mathbf{f}^{\mathrm{T}}(S_{l,i}, P_{l,i})\mathbf{w}_{l,i})\right] \mathbf{f}(S_{l,i}, P_{l,i}). \tag{31}$$

If the condition in (30) is not fulfilled, the weights are not updated in the time interval and $\mathbf{w}_{l,i+1} = \mathbf{w}_{l,i}$.

V. Simulation Results

In this section, we present numerical results for the evaluation of the proposed multi-agent reinforcement learning algorithm. In the figures, we refer to our proposed algorithm as proposed MARL. For the simulations, we consider $T = 1000$ independent random channel and energy realizations, where each realization corresponds to an episode in which the EH nodes harvest energy I times. The duration τ of each time interval is assumed to be equal to one time unit.

The amounts of energy $E_{l,i}$, $l \in \{1, 2\}$, harvested by N_1 and N_2 in time interval i are taken from a uniform distribution with maximum value $E_{\max,l}$. The channel coefficients $h_{l,i}$ are taken from an i.i.d. Rayleigh fading process with zero mean and unit variance. Additionally, a bandwidth $W = 1$MHz is assumed to be available for the transmission from N_1 to N_2, and from N_2 to N_3. The noise variance is set to $\sigma^2 = 1$ and a tolerable quantization error of one percent is considered. The step size δ used in the definition of the action set $\mathcal{A}_l$ is set to $\delta = 0.02B_{\max,l}$. The learning rate α and the ϵ parameter used in the ϵ-greedy policy are reduced in each time instant and are defined as $\alpha = 1/i$ and $\epsilon = 1/i$, respectively. Furthermore, the discount factor γ is selected as $\gamma = 0.9$. To compare the performance of our proposed MARL algorithm, we consider the following alternative approaches:

- Offline optimum: In this case, the offline optimum policy is obtained by assuming that a central entity has perfect non-causal knowledge regarding the EH processes, the data arrival processes and the channel fading processes of nodes N_1 and N_2. With this knowledge, an optimization problem is formulated at the central entity to maximize the throughput in (1) considering the constraints in (2), (3), (4) and (5) in order to jointly find the transmit powers $p_{1,i}$ and $p_{2,i}$ of nodes N_1 and N_2, respectively.
- Learning - No Cooperation [30]: As described in [30], in this approach it is assumed that the EH nodes have only causal knowledge regarding their own EH processes, data arrival processes and channel fading processes. No cooperation between the EH nodes to exchange their causal knowledge is assumed and the transmission policy is obtained by solving independent reinforcement learning problems at nodes N_1 and N_2.
- Hasty policy: This policy is obtained by considering that the nodes have only causal knowledge regarding their own EH processes, data arrival processes and channel fading processes. It is a simple approach that consists of depleting the battery of node N_1 in each time interval to transmit the maximum possible amount of data to node N_2. At node N_2, the hasty policy aims at depleting the data buffer by selecting the maximum transmit power value in each time interval that fulfils the data causality constraint in (4).

In Fig. 3, we compare the sum of the throughput over the I time intervals, measured in bits, for different values of the fraction τ_{sig}/τ of the duration of the time interval assigned to the signaling phase. For this simulation, $I = 100$ time intervals are considered. The battery sizes of

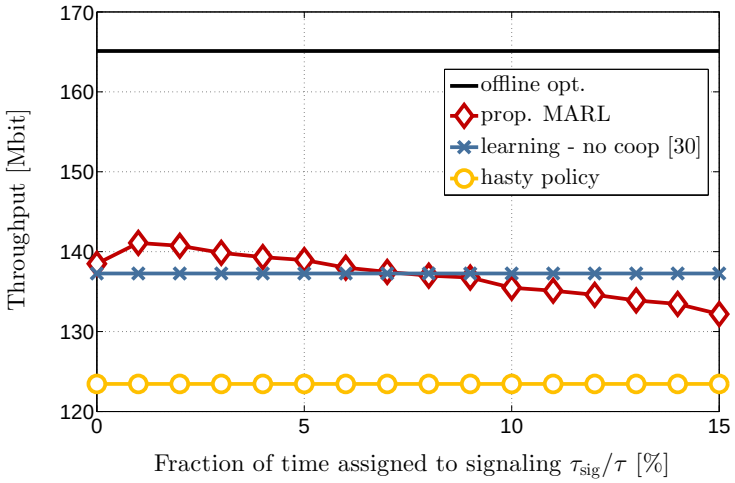

Fig. 3: Average throughput versus fraction of time assigned to signaling τ_{sig}/τ.

the nodes are assumed to be $B_{\max,1} = B_{\max,2} = 2E_{\max}$, where $E_{\max}/2\sigma^2 = 5$dB. Furthermore, to focus on the achieved throughput when the data arrival process at node N_1 is not the limiting factor, it is assumed that node N_1 has a data buffer of infinite size and has always data available for transmission. At node N_2, the data buffer size $D_{\max,2}$ is calculated as the amount of data node N_2 would receive in one time interval if the transmit power would be $p_{1,i} = B_{\max,1}/\tau$ and the channel coefficient would be $|h_{1,i}| = 1$, i.e., $D_{\max,2} = W\tau \log_2(1 + B_{\max,1}/\tau)$. With this parameters, the numbers of bits required to transmit the signaling are $L_{E_{l,i}} = 9$ bits, $L_{B_{l,i}} = 10$ bits, $L_{D_{l,i}} = 28$ bits and $L_{E_{l,i}} = 7$ bits. As expected, the largest throughput is achieved by the offline optimum policy at the cost of requiring perfect non-causal knowledge regarding the EH processes and the channel fading processes of nodes N_1 and N_2. Furthermore, the throughput achieved by the learning approach without cooperation and the hasty policy is flat because they do not consider a signaling phase. Therefore, the complete duration of the time interval is used for the transmission of data. The achieved throughput of the proposed MARL depends on the time assigned for the signaling. For $\tau_{\text{sig}}/\tau < 7\%$, the proposed MARL outperforms the other two approaches which also consider only causal knowledge. The reason for this improvement is that by including the signaling phase, nodes N_1 and N_2 exchange their causal knowledge and are able to learn a transmission policy that adapts to the battery levels, data buffer levels and channel coefficients of both nodes. The maximum throughput of the proposed MARL is achieved at

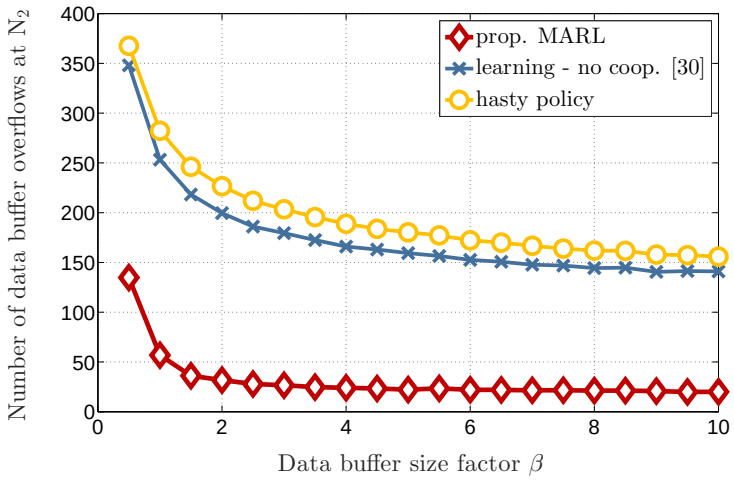

Fig. 4: Average throughput versus $E_{\max}/(2\sigma^2)$.

$\tau_{\text{sig}}/\tau = 1\%$. For $\tau_{\text{sig}}/\tau < 1\%$, the throughput is reduced because, as shown in (22), the relation between τ_{sig} and the transmit power $p_{\text{sig},l,i}$ required to transmit the signaling is not linear and the smaller τ_{sig}, the over-proportionally larger $p_{\text{sig},l,i}$. As $p_{\text{sig},l,i}$ increases, the probability of not having enough energy in the battery to fulfil this requirement increases. Consequently, the nodes do not have enough energy to transmit during the signaling phase and to exchange their causal knowledge. When τ_{sig}/τ increases to values beyond 1%, the achieved throughput decreases. Even though in this case of $\tau_{\text{sig}}/\tau > 1\%$, the EH nodes have a longer signaling phase to exchange their causal knowledge, can therefore use less power for the transmission of the signaling and consequently, can save energy for data transmission, less time is left for the transmission of data. As a result, the power required to transmit a certain amount of data increases.

The number of data buffer overflow situations at node N_2 versus the data buffer size is shown in Fig. 4 for $I = 1000$. As in Fig. 3, it is assumed that node N_1 has always data to transmit, and the parameters are chosen as $E_{\max}/2\sigma^2 = 5$dB and $B_{\max,1} = B_{\max,2} = 2E_{\max}$. To vary the data buffer size at node N_2, the factor β is introduced and $D_{\max,2}$ is calculated as $D_{\max,2} = W\tau \log_2(1 + \beta B_{\max,1}/\tau)$. For the proposed MARL, $\tau_{\text{sig}}/\tau = 1\%$ is selected. We have omitted the result of the offline optimum because due to (5), data buffer overflow situations are not allowed in the solution of the optimization problem. It can be seen that, as the data buffer size increases, the number of data buffer overflow situations is reduced for all the approaches. This is

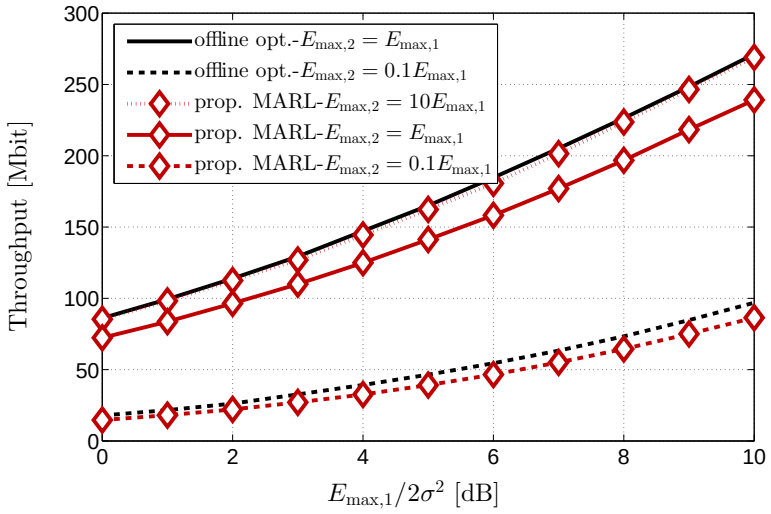

Fig. 5: Average throughput versus data buffer size factor β.

because the probability of having data buffer overflow situations decreases with increasing data buffer size. For $\beta = 0.5$, the proposed MARL has 2.6 times less data buffer overflow situations than the learning approach without cooperation and 2.7 times less than the hasty policy. The relative difference between the proposed MARL and the two other approaches increases for larger data buffer sizes. For example, for $\beta = 10$, the proposed MARL has 7 times less data buffer overflow situations than the learning approach without cooperation and 8 times less than the hasty policy. The better performance of the proposed MARL results from the fact that by exchanging the causal knowledge during the signaling phase, node N_1 knows the data buffer level of node N_2 and can limit the amount of transmitted data when the data buffer of node N_2 is almost full. It should be noted that although the proposed MARL is able to significantly reduce the number of data buffer overflow situations, it cannot reduce it to zero. This is because non-causal knowledge would be required in order to adapt the transmission policy according to the amounts of energy that will be harvested in the future as well as the future channel coefficients.

In Fig. 5, we compare the performance of the offline optimum policy and the proposed MARL for several values of $E_{\max,1}$. $I = 100$ time intervals are considered and for the proposed MARL, $\tau_{\text{sig}}/\tau = 1\%$ is selected. We additionally evaluate the effect of the maximum amount of energy which nodes N_1 and N_2 can harvest. For this purpose, we consider three different cases, i.e., $E_{\max,2} = 10E_{\max,1}$, $E_{\max,2} = E_{\max,1}$ and $E_{\max,2} = 0.1E_{\max,1}$. For the first case, i.e. $E_{\max,2} = 10E_{\max,1}$, the offline optimum policy cannot be applied because battery overflow situations cannot

be avoided at node N_2 when it harvests much more energy than node N_1. This is due to the fact that node N_2 has more energy available in its battery than what is needed to retransmit the data it receives from node N_1. To allow battery overflow situations at node N_2, a different optimization problem would need to be considered which is out of the scope of our work. In all the three cases, the throughput increases when the amount of harvested energy increases. The largest throughput is achieved by the proposed MARL for the case when $E_{\text{max},2} = 10E_{\text{max},1}$ and this throughput is close to the offline optimum performance for $E_{\text{max},2} = E_{\text{max},1}$. This is because harvesting more energy at node N_2 cannot lead to a larger throughput if the amount of harvested energy is not increased at N_1. The throughput is limited by the amount of data node N_1 can transmit which in turn is limited by the amount of energy node N_1 harvests, which for the two cases, $E_{\text{max},2} = 10E_{\text{max},1}$ and $E_{\text{max},2} = E_{\text{max},1}$, is in a similar order of magnitude. For $E_{\text{max},2} = E_{\text{max},1}$, the performance of the proposed MARL is reduced compared to the case when $E_{\text{max},2} = 10E_{\text{max},1}$. This is because there is less energy available at N_2. As a result, in each time interval, N_2 allocates less energy for data transmission. For the case when $E_{\text{max},2} = 0.1E_{\text{max},1}$, the performance of the proposed MARL is close to the performance of the offline optimum policy. This is due to the fact that in this case, node N_2 is the bottleneck because it harvests on average much less energy than node N_1. Both approaches, the offline optimum policy and the proposed MARL, limit the amount of data node N_1 transmits while aiming at maximizing the achieved throughput in each time interval.

Fig. 6 and Fig. 7 show the effect of having a data arrival process at node N_1. For this simulation, we consider that the number of data packets arriving in each time interval follows a Poisson distribution where the average number of data packets arriving per time interval is given by λ. $I = 1000$ time intervals are considered. Moreover, we consider a packet size of 200 kbit and 500 kbit for Fig. 6 and Fig. 7, respectively. Additionally, $E_{\text{max},1} = E_{\text{max},2} = E_{\text{max}}$, $E_{\text{max}}/2\sigma^2 = 5$dB and $B_{\text{max},1} = B_{\text{max},2} = 2E_{\text{max}}$ hold. The data buffer sizes are calculated as $D_{\text{max},1} = D_{\text{max},2} = W\tau \log_2(1+2E_{\text{max}}/\tau)$. The offline optimum policy is not considered because the feasibility of the optimization problem depends on each particular realization of the data arrival process. In Fig. 6, it can be seen that for $\lambda < 5$, the proposed MARL, the learning approach without cooperation, and the hasty policy achieve almost the same performance. This is because for $\lambda < 5$, the data buffer is almost empty all the time. Therefore, data buffer overflow situations are unlikely and the data packets received by node N_1 can be retransmitted by node N_2

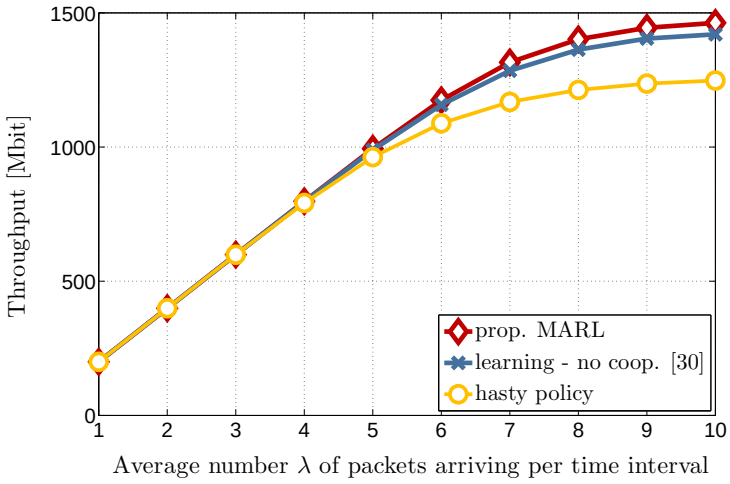

Fig. 6: Average throughput versus average number λ of packets arriving per time interval. Packet size of 200 kbit.

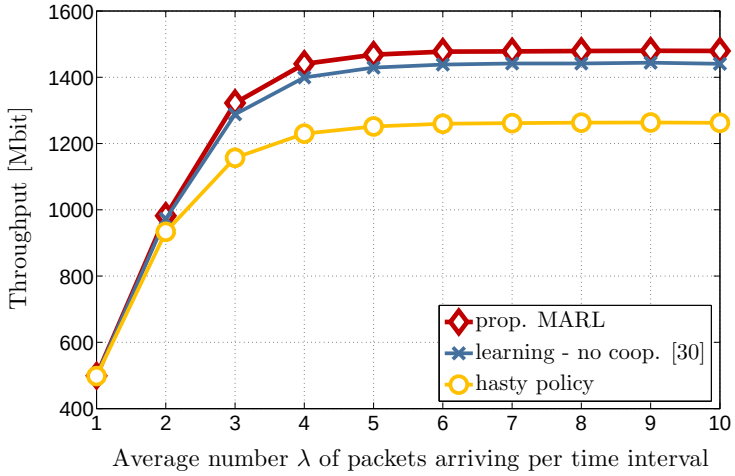

Fig. 7: Average throughput versus average number λ of packets arriving per time interval. Packet size of 500 kbit.

to node N_3. For example, for $\lambda = 2$, two data packets are received on average per time interval and the total achieved throughput adding up all the time intervals is $(2\lambda I) = 400$ Mbit. As the number of data packets received per time interval increases, the performance of the learning approach without cooperation and of the hasty policy decreases because these two approaches do nothing to prevent data buffer overflow situations at node N_2. From Fig. 7, it can be seen that for a packet size of 500 kbit, the performance of the three approaches saturates at approximately

$\lambda = 5$. For $\lambda \geq 5$, data buffer overflow situations at node N_1 cannot be avoided because the number of data packets received per time interval is larger than the data buffer capacity.

Finally, in Fig. 8, we evaluate the convergence of the proposed MARL. For this purpose, we compare the normalized throughput, i.e. the number of bits transmitted divided by the number I of time intervals versus the number I of time intervals. In addition to the proposed MARL and the learning approach without cooperation, we evaluate the performance of the proposed feature functions by implementing the proposed MARL using two standard approximation techniques, namely, fixed sparse representation (FSR) and radial basis functions (RBF) [37]. Both, FSR and RBF are low-complexity techniques used to represent the continuous states. For each EH node N_l, $l \in \{1, 2\}$, the state S_i, observed after the signaling phase, lies in an 8-dimensional space given by $E_{l,i}$, $B_{l,i}$, $h_{l,i}$ and $D_{l,i}$. In FSR, each dimension is split in tiles and a binary feature function is assigned to each tile. A given feature function is equal to one if the corresponding variable is in the tile and zero otherwise [37]. In our implementation, the tiles are generated by quantizing each dimension using the step size δ used in the definition of the action spaces $\mathcal{A}_l$. In RBF, each feature function has a Gaussian shape that depends on the distance between a given state and the center of the feature [33], [37]. In contrast to FSR, in RBF a given state is represented by more than one feature function. In Fig. 8, it can be seen that the proposed MARL and the learning approach without cooperation converge at approximately the same number of iterations. This is due to the fact that both approaches are based on the SARSA update. However, since the proposed MARL considers the full cooperation among the EH nodes to exchange their causal knowledge, it can achieve a larger throughput. The number of feature functions required by a learning approach affects the learning rate and the performance. Consequently, our proposed MARL outperforms FSR and RBF because they require a larger number of feature functions compared to the proposed MARL which only needs six.

To summarize the simulation results, it can be seen that with a proper selection of τ_{sig}, the proposed MARL, which considers cooperation between the EH nodes, outperforms other approaches without cooperation between the nodes and which also only consider causal knowledge. This means, that reserving a fraction of time for the exchange of signaling among the nodes is more beneficial than assuming no cooperation at all, even though the time dedicated to data transmission is reduced in order to include the signaling phase. When the nodes cooperate with each other, a higher throughput can be achieved. Additionally, the proposed MARL reduces the

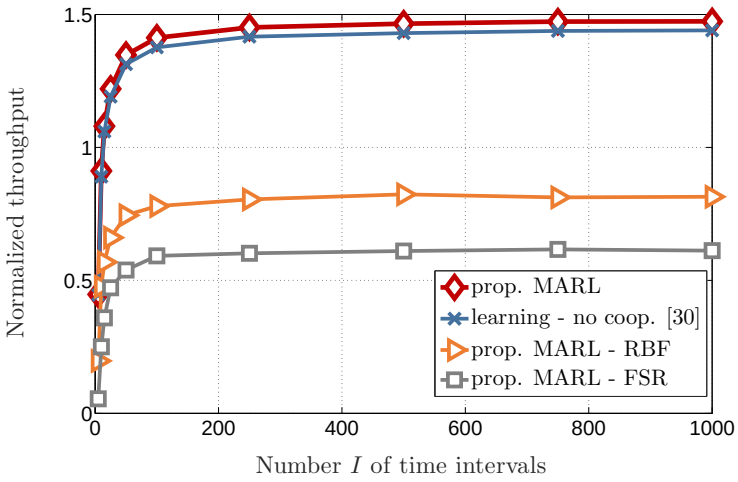

Fig. 8: Normalized throughput versus number I of time intervals.

number of data buffer overflows at N_2 as compared to the other approaches. This implies a reduction in the number of required retransmissions.

VI. Conclusion

We have investigated an EH two-hop communication scenario where only causal knowledge regarding the EH processes, the data arrival processes and the channel fading processes was assumed at the EH transmitter and at the EH relay. We considered the case when a signaling phase is available in each time interval. This signaling phase is used by the EH nodes to cooperate with each other by exchanging their causal knowledge. After the signaling phase, the EH nodes exploit the obtained knowledge to find transmission policies which adapt to the battery levels, data buffer levels and channel coefficients of the EH nodes and which aim at maximizing the throughput. We modelled the problem as a Markov game and proposed a multi-agent reinforcement learning algorithm to find the transmission policies at the transmitter and at the relay. Furthermore, we have provided convergence guarantees for the proposed algorithm. Through several simulation results we have shown that a larger throughput can be achieved when cooperation among the EH nodes is considered, compared to the case when no cooperation is assumed even after the overhead is subtracted from the number of bits transmitted. Moreover, we have shown the trade-off between the duration of the signaling phase and the performance of the proposed algorithm and we have shown that the number of data buffer overflow situations is reduced when our

proposed algorithm is considered. The distributed nature of our proposed algorithm makes it suitable for more complex relay networks, e.g., multi-hop networks.

REFERENCES

[1] W. Dargie and C. Poellabauer, Fundamentals of Wireless Sensor Networks: Theory and Practice. John Wiley and Sons, 2010.

[2] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang, "Energy harvesting wireless communication: A review of recent advances," IEEE J. Sel. Areas Commun., vol. 33, no. 3, pp. 360–381, March 2015.

[3] M.-L. Ku, W. Li, Y. Chen, and K. J. R. Liu, "Advances in energy harvesting communications: Past, present, and future challenges," IEEE Commun. Surveys Tutorials, vol. 18, no. 2, pp. 1384–1412, 2nd Quart. 2016.

[4] D. Gündüz, A. Yener, A. Goldsmith, and H. V. Poor, "The multi-way relay channel," IEEE Trans. Inform. Theory, vol. 59, no. 1, pp. 51–63, January 2013.

[5] E. Yilmaz, R. Zakhour, D. Gesbert, and R. Knopp, "Multi-pair two-way relay channel with multiple antenna relay station," in Proc. IEEE Int. Conf. Commun. (ICC), Cape Town, May 2010, pp. 1–5.

[6] K. Tutuncuoglu and A. Yener, "Optimum transmission policies for battery limited energy harvesting nodes," IEEE Trans. Wireless Commun., vol. 11, no. 3, pp. 1180–1189, March 2012.

[7] O. Ozel, K. Tutuncuoglu, J. Yang, S. Ulukus, and A. Yener, "Transmission with energy harvesting nodes in fading wireless channels: Optimal policies," IEEE J. Sel. Areas Commun., vol. 29, no. 8, pp. 1732–1743, September 2011.

[8] O. Orhan and E. Erkip, "Throughput maximization for energy harvesting two-hop networks," in Proc. IEEE Int. Symp. Inform. Theory (ISIT), Istanbul, July 2013, pp. 1596–1600.

[9] ——, "Optimal transmission policies for energy harvesting two-hop networks," in Proc. Annual Conf. Inform. Sciences Syst. (CISS), Princeton, March 2012, pp. 1–6.

[10] X. Li, A. Ortiz, H. Al-Shatri, A. Klein, and T. Weber, "Delay-constrained data transmission for energy harvesting transmitter," in Proc. Int. ITG Conf. Syst. Commun. and Coding (SCC), Hamburg, February 2017, pp. 1–6, accepted.

[11] O. Ozel, J. Yang, and S. Ulukus, "Optimal broadcast scheduling for an energy harvesting rechargeable transmitter with a finite capacity battery," IEEE Trans. Wireless Commun., vol. 11, no. 6, pp. 2193–2203, June 2012.

[12] M. Fu, A. Arafa, S. Ulukus, and W. Chen, "Delay minimal policies in energy harvesting broadcast channels," in Proc. IEEE Int. Conf. Commun. (ICC), Kuala Lumpur, May 2016, pp. 1–6.

[13] J. Yang and S. Ulukus, "Optimal packet scheduling in a multiple access channel with energy harvesting transmitters," J. Commun. and Networks, vol. 14, no. 2, pp. 140–150, April 2012.

[14] Z. Wang, V. Aggarwal, and X. Wang, "Iterative dynamic water-filling for fading multiple-access channel with energy harvesting," IEEE J. Sel. Areas Commun., vol. 33, no. 3, pp. 382–395, March 2015.

[15] D. Gündüz and B. Devillers, "Two-hop communication with energy harvesting," in Proc. IEEE Int. Workshop Comput. Advances Multi-Sensor Adaptive Process. (CAMSAP), San Juan, December 2011, pp. 201–204.

[16] O. Orhan and E. Erkip, "Energy harvesting two-hop communication networks," IEEE J. Sel. Areas Commun., vol. 33, no. 12, pp. 2658–2670, December 2015.

[17] A. Ortiz, H. Al-Shatri, X. Li, T. Weber, and A. Klein, "Throughput maximization in two-hop energy harvesting communications," in Proc. Int. Symp. Wireless Commun. Syst. (ISWCS), Brussels, August 2015, pp. 291–295.

[18] Y. Luo, J. Zhang, and K. B. Letaief, "Optimal scheduling and power allocation for two-hop energy harvesting communication systems," IEEE Trans. Wireless Commun., vol. 12, no. 9, pp. 4729–4741, September 2013.

[19] B. Varan and A. Yener, "Two-hop networks with energy harvesting: The (non-)impact of buffer size," in Proc. IEEE Global Conf. Signal Inform. Process. (GlobalSIP), Austin, December 2013, pp. 399–408.

[20] K. Tutuncuoglu, B. Varan, and A. Yener, "Optimum transmission policies for energy harvesting two-way relay channels," in Proc. IEEE ICC'13 Workshop on GBA: Energy Efficient Wireless and Wired Network Solutions, Budapest, June 2013, pp. 586–590.

[21] B. Varan and A. Yener, "The energy harvesting two-way decode-and-forward relay channel with stochastic data arrivals," in Proc. IEEE Global Conf. Signal Inform. Process. (GlobalSIP), Austin, December 2013, pp. 371–374.

[22] J. Lei, R. Yates, and L. Greenstein, "A generic model for optimizing single-hop transmission policy of replenishable sensors," IEEE Trans. Wireless Commun., vol. 8, no. 2, pp. 547–551, February 2009.

[23] S. Luo, R. Zhang, and T. J. Lim, "Optimal save-then-transmit protocol for energy harvesting wireless transmitters," IEEE Trans. Wireless Commun., vol. 12, no. 3, pp. 1196–1207, March 2013.

[24] W. Zeng, Y. R. Zheng, and R. Schober, "Online resource allocation for energy harvesting downlink multiuser systems: Precoding with modulation, coding rate, and subchannel selection," IEEE Trans. Wireless Commun., vol. 14, no. 10, pp. 5780–5794, October 2015.

[25] A. Minasian, S. ShahbazPanahi, and R. S. Adve, "Energy harvesting cooperative communication systems," IEEE Trans. Wireless Commun., vol. 13, no. 11, pp. 6118–6131, November 2014.

[26] F. Amirnavaei and M. Dong, "Online power control for cooperative relaying with energy harvesting," in Proc. Asilomar Conf. Signals, Syst. Computers, Pacific Grove, November 2015, pp. 817–822.

[27] P. Blasco, D. Gündüz, and M. Dohler, "A learning theoretic approach to energy harvesting communication system optimization," IEEE Trans. Wireless Commun., vol. 12, no. 4, pp. 1872–1882, April 2013.

[28] A. Ortiz, H. Al-Shatri, X. Li, T. Weber, and A. Klein, "Reinforcement learning for energy harvesting point-to-point communications," in Proc. IEEE Int. Conf. Commun. (ICC), Kuala Lumpur, May 2016, pp. 1–6.

[29] P. Blasco and D. Gündüz, "Multi-access communications with energy harvesting: A multi-armed bandit model and the optimality of the myopic policy," IEEE J. Sel. Areas Commun., vol. 33, no. 3, pp. 585–597, March 2015.

[30] A. Ortiz, H. Al-Shatri, X. Li, T. Weber, and A. Klein, "A learning based solution for energy harvesting decode-and-forward two-hop communications," in Proc. IEEE Global Commun. Conf. (Globecom), Washington, December 2016, pp. 1–7.

[31] M. L. Littman, "Markov games as a framework for multi-agent reinforcement learning," in Proc. Int. Conf. Machine Learning, New Brunswick, July 1994, pp. 157–163.

[32] ——, "Value-function reinforcement learning in Markov games," J. Cognitive Syst. Research, vol. 2, no. 1, pp. 55–66, October 2001.

[33] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction. MIT Press, 1998.

[34] M. Lauer and M. Riedmiller, "An algorithm for distributed reinforcement learning in cooperative multi-agent systems," in Proc. Int. Conf. Machine Learning, Stanford, June 2000, pp. 535–542.

[35] S. Russell and P. Norvig, Artificial Intelligence: A Modern Approach, 3rd ed. Prentice Hall, 2010.

[36] G. J. Gordon, "Reinforcement learning with function approximation converges to a region," Advances Neural Inform. Process. Syst., pp. 1040–1046, 2001.

[37] A. Geramifard, T. J. Walsh, S. Tellex, G. Chowdhary, N. Roy, and J. P. How, "A tutorial on linear function approximators for dynamic programming and reinforcement learning," Foundations and Trends in Machine Learning, vol. 6, no. 4, pp. 375–454, December 2013.